



# Demonstration of Cost-Effective, High-Performance Computing at Performance and Reliability Levels Equivalent to a 1994 Vector Supercomputer

Theresa Babrauckas  
Glenn Research Center, Cleveland, Ohio

## The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the Lead Center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.

- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized data bases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at <http://www.sti.nasa.gov>
- E-mail your question via the Internet to [help@sti.nasa.gov](mailto:help@sti.nasa.gov)
- Fax your question to the NASA Access Help Desk at (301) 621-0134
- Telephone the NASA Access Help Desk at (301) 621-0390
- Write to:  
NASA Access Help Desk  
NASA Center for Aerospace Information  
7121 Standard Drive  
Hanover, MD 21076



# Demonstration of Cost-Effective, High-Performance Computing at Performance and Reliability Levels Equivalent to a 1994 Vector Supercomputer

Theresa Babrauckas  
Glenn Research Center, Cleveland, Ohio

National Aeronautics and  
Space Administration

Glenn Research Center

## Acknowledgments

The author wishes to express her gratitude to the following people for their help in writing this report: Glenn Genzlinger and Dan Minior, Pratt & Whitney; and Leigh Ann Tanner, NASA Ames Research Center.

Available from

NASA Center for Aerospace Information  
7121 Standard Drive  
Hanover, MD 21076  
Price Code: A03

National Technical Information Service  
5285 Port Royal Road  
Springfield, VA 22100  
Price Code: A03

# **DEMONSTRATION OF COST-EFFECTIVE, HIGH-PERFORMANCE COMPUTING AT PERFORMANCE AND RELIABILITY LEVELS EQUIVALENT TO A 1994 VECTOR SUPERCOMPUTER**

Theresa Babrauckas  
National Aeronautics Space Administration  
Glenn Research Center  
Cleveland, Ohio 44135

## **SUMMARY**

The Affordable High Performance Computing (AHPC) project demonstrated that high-performance computing based on a distributed network of computer workstations is a cost-effective alternative to vector supercomputers for running CPU and memory intensive design and analysis tools. The AHPC project created an integrated system called a Network Supercomputer. By connecting computer workstations through a network and utilizing the workstations when they are idle, the resulting distributed-workstation environment has the same performance and reliability levels as the Cray C90 vector supercomputer at less than 25 percent of the C90 cost. In fact, the cost comparison between a Cray C90 Supercomputer and Sun workstations showed that the number of distributed networked workstations equivalent to a C90 costs approximately 8 percent of the C90.

## **INTRODUCTION**

The Affordable High Performance Computing (AHPC) project was a Computational Aerosciences (CAS) project within the High Performance Computing and Communications Program. The goals of the project were to demonstrate end-to-end reductions in cost and time, find a solution for aerospace design applications on heterogeneous systems, and demonstrate the affordability of high-performance computing using a distributed network of workstations. This paper contains the analysis of the CAS milestone demonstrating cost-effective, high-performance computing at performance and reliability levels equivalent to 1994 vector supercomputers at 25 percent of the capital cost.

The specific CAS milestones for the AHPC project are listed below.

- Release AHPC Cooperative Agreement Notice (CAN) Proposal (October 1994)
- Sponsor AHPC CAN Proposal Conference (January 1995)
- AHPC CAN Proposals due (February 1995)
- Award AHPC CAN (May 1995)
- Demonstrate price/performance (September 1997)

The first four milestones were completed on time. The AHPC CAN was awarded to Pratt & Whitney and its partners on May 31, 1995. The final milestone of the price and performance demonstration is complete. This report outlines the details of that demonstration. The tables and charts illustrate the performance, reliability, and cost comparisons in detail.

## **PERFORMANCE**

Table I contains data that shows the number of workstations needed to perform the same as one node of the Cray C90. A small case of a CFD code was run serially on the Cray C90 at NASA Ames Research Center. The code was run for many years on a Cray. The same code was parallelized and ported to Sun workstations. The only additional changes made were for getting clean compiles and executions. The code was not optimized for Sun workstations. The case was run in accordance with the AHPC project on two different Sun workstation models at Pratt & Whitney, both serially and in parallel. Several different parallel breakups were run on the Sun network with 2, 4, 6, 8, and 10 workstations. All codes were run on dual processor machines except those that were run on the C90.

The following definitions apply to table I:

For the Sun workstation cases, “node” is equivalent to “workstation processor.” Required speedup is equivalent to (wall time of Sun case)/(wall time of C90 case), which indicates the performance ratio between the Sun cases and C90 case. For example, the six-node Ultra 2 case runs 1.08 times slower than the case on one node of the C90.

TABLE I.—CRAY C90 SUPERCOMPUTER VERSUS NETWORK OF WORKSTATIONS

Platform	Mode	Nodes	Wall time, sec	Required speedup
C90	Serial	1	129.33	1.00
Sun SPARCstation 20	Serial	1	1362.00	10.53
Model 612 (S20M612)				
S20M612	Parallel	2	796.50	6.16
S20M612	Parallel	4	504.88	3.90
S20M612	Parallel	6	394.04	3.05
S20M612	Parallel	8	343.25	2.65
S20M612	Parallel	10	298.07	2.30
Sun Ultra 2 M2200	Serial	1	547.50	4.22
Sun Ultra 2 M2200	Parallel	2	302.83	2.33
Sun Ultra 2 M2200	Parallel	4	201.13	1.55
Sun Ultra 2 M2200	Parallel	6	140.50	1.08
Sun Ultra 2 M2200	Parallel	8	123.00	0.95
Sun Ultra 2 M2200	Parallel	10	109.97	0.85

It is important to note that all workstation cases were run during normal business hours on a nondedicated, open network of desktop machines with either a 10-MB switched Ethernet (Sun SPARCstation) or a 100-MB switched Ethernet (Sun Ultra 2). The 100 MB switched Ethernet is the recommended choice as a result of the ATM network study performed under this project (ref. 1).

In addition, each desktop Sun workstation has two processors. One processor is available to users on the workstation network for running parallel jobs. The cost analysis below includes the cost of the whole machine even though only one processor was used by the parallel benchmark.

Taking the above into consideration, table I shows that eight Sun Ultra 2 Model 2200’s are equivalent to a single node of the Cray C90 (123.00 versus 129.33 sec). Table I also shows that the Cray C90 (129.33 sec) runs 2.3 times as fast as the 10-node SPARCstation 20 (298.07 sec).

## RELIABILITY

Even if it can be shown that a network of workstations performs the same as one node of the C90, the reliability levels of both platforms need to be equivalent in order to achieve a fair cost comparison. Two types of reliability (or availability) metrics are tracked regularly on both the C90 and the Pratt & Whitney Sun workstation network and are described below. In addition, Pratt & Whitney tracks parallel job reliability.

Scheduled availability is the amount of time that a computer is available outside of regularly scheduled maintenance. The percentage availability is the ratio of uptime per specific time frame. For the Cray C90, there are 80 scheduled hours of downtime each year that are devoted to normal maintenance. The percentage availability is

$$\begin{aligned}
 \text{Scheduled availability} &= (1 \text{ year} - 80 \text{ hr}) / (1 \text{ year}) \\
 &= (365 \text{ days} \times 24 \text{ hr} - 80 \text{ hr}) / (365 \text{ days} \times 24 \text{ hr}) \\
 &= (8760 \text{ hr} - 80 \text{ hr}) / (8760 \text{ hr}) \\
 &= 99.1 \text{ percent availability}
 \end{aligned}$$

At Pratt & Whitney in East Hartford, scheduled workstations maintenance averages 1 hr per week. The effect is that parallel computing cannot be performed on average, 1 hr per week. The scheduled availability for the distributed workstations is

$$\begin{aligned} \text{Scheduled availability} &= (1 \text{ week} - 1 \text{ hr}) / (1 \text{ week}) \\ &= (7 \text{ days} \times 24 \text{ hr} - 1 \text{ hr}) / (7 \text{ days} \times 24 \text{ hr}) \\ &= (168 \text{ hr} - 1 \text{ hr}) / (168 \text{ hr}) \\ &= 99.4 \text{ percent availability} \end{aligned}$$

Gross availability is the fraction of time that a computer is available regardless of whether it is scheduled to be up or down. The gross availability of the C90 at NASA Ames Research Center in a 6-month period (from March 17, 1997 to September 1, 1997) is shown in the following graph:





Figure 1.—Gross availability of C90.

The average gross availability for the C90 during this 6-month period was 99.0 percent.

The gross availability of Sun workstations for a recent 38-month period at Pratt & Whitney is shown below. The triangular data is the gross availability for each month. The rectangular data is the average gross availability for the previous months.



The average gross availability for the 38-month period was 99.1 percent. The most recent gross availability for the 38-month period is 99.7 percent. If an individual parallel job requires eight workstations, the predicted gross availability for the set of workstations would be 99.7<sup>8</sup> or 97.6 percent.

On the basis of the data received from Pratt & Whitney, we have concluded that the Pratt & Whitney Sun workstation network has scheduled and gross availability rates that are equivalent to the scheduled and gross availability levels of the Cray C90. Also, parallel job completion rates on the Pratt & Whitney Sun workstation network are essentially equivalent to the Cray C90 assuming that the C90 has a 100 percent job completion rate. Therefore, parallel computing on a network of workstations can be as reliable as running serially on a supercomputer.

Table II outlines the costs for the number of equivalent Sun Ultra 2 workstations as determined in table I. From the performance data, it can be concluded that eight Sun Ultra 2 workstations match the performance of a single C90 processor. Consequently, 128 Sun Ultra 2 workstations would match the capacity of a whole 16-processor C90. With the given failure rate of some of the Sun workstations in the open network, additional "fail-over" machines need to be included. If it is estimated that an additional 2 workstations are needed, then a total of 10 Sun Ultra 2's would be equivalent in performance and reliability to 1 C90 node. Table II also includes the cost of the number of Sun SPARCstation workstations needed to achieve approximately half of the performance of one node of the C90. The Scientific & Engineering Workstation Procurement (SEWP) government contract was used to determine the cost of the workstations, file servers, and yearly maintenance prices. The networking and system administration costs are typical costs incurred at NASA Glenn Research Center. The costs assumed for table II are as follows:



- Unix workstation cost = \$19 K each for the SPARCstation (dual processor, 128-MB memory, 1-GB hard disk), and \$22 K each for the Ultra 2 (dual processor, 200-MHz chip, 192-MB memory, 1-GB hard disk)
- Unix workstation network cost = \$1144 per workstation (includes cost of equipment and Full Time Equivalent (FTE) support)
- Unix workstation server cost = \$30 K each
- Need one server for every 80 workstations
- Unix workstation maintenance: \$1500 per year per workstation
- Unix workstation system administration (SA) costs = 1 system administrator for 50 Unix workstations. Cost of one SA = \$75 K.
- Cost of single node of the Cray C90 is arrived at from the cost (\$48.5 M) of the 16-node Cray C90

TABLE II.—NETWORK OF WORKSTATIONS COST VERSUS CRAY C90 COST

Platform	Number of nodes	Hardware costs	Network costs	Server costs	Maintenance costs	System administration costs	Total costs	Ratio of 10 workstations to 1 C90 node, percent
Sun SPARCstation 20M612 (1995)	10	\$190,000	\$11,440	\$3,750	\$15,000	\$15,000	\$235,190	6.78
Sun Ultra 2 (1996)	10	220,000	11,440	3,750	15,000	15,000	265,190	7.64
Cray C90 (1994)	1	3,031,250	Included in hardware costs	N/A	300,000	137,500	3,468,750	
Cray C90 (1994)	16	48,500,000	Included in hardware costs	N/A	4,800,000	2,200,000	55,500,000	

The above table shows that a network of workstations equivalent to a single node of the Cray C90 super-computer costs no more than 8 percent of the capital cost of the C90. This percentage clearly exceeds the milestone goal of 25 percent.

One can argue that comparing 1994 and 1996 technology is not a fair analysis. In response to this argument, a cost comparison between the Cray T3E (1996 technology) and Sun Ultra 2 workstations (also 1996 technology) is included in the addendum at the end of this paper. However, it is traditional to include well-known computer hardware in benchmark comparison. Also with the rapidity of new workstation hardware releases, increased performance at reduced costs will continue.

## CONCLUSION

As shown in the performance, reliability and cost data above, the CAS Level 1 milestone, “Demonstrate Cost-effective, High-performance Computing at Performance and Reliability Levels Equivalent to 1994 Vector Supercomputers at 25 percent of the Capital Cost” was successfully completed. Distributed networks of computers such as the Sun workstation network at Pratt & Whitney are a viable and affordable alternative to the traditional supercomputer.

This study does not advocate replacing all Cray Supercomputers with dedicated computer clusters. Instead it advocates networking together the existing desktop workstations within an organization. These workstations may have been purchased for other purposes. A significant number of workstations are underutilized most of the time. Research laboratories and other U.S. aeroengine companies besides Pratt & Whitney may have supercomputer class power sitting idle. This study, and the Affordable High Performance Computing project in general, serves to make organizations aware of ways to tap their idle resources.

## APPENDIX—COST COMPARISON BETWEEN CRAY T3E AND SUN ULTRA 2 WORKSTATIONS

An argument may be made about the appropriateness of comparing the cost of the Cray C90 to the Sun Ultra 2 workstations. The Cray C90 is 1994 computer technology while the Sun Ultra 2 is 1996 computer technology. Some may argue that this comparison is not relevant or accurate since the computer technology compared has a difference of 2 years. The Cray T3E is 1996 computer technology, which is closer to the Sun Ultra 2 workstation's computer technology than the Cray C90. Therefore, it is valuable to compare the cost, performance, and reliability levels of the Cray T3E and the Sun Ultra 2 workstation. The reader is cautioned that the comparison to follow is more of a "back of the envelope" comparison since the time and resources were not available to devote to a full analysis like the Cray C90 and Sun Ultra 2 workstation comparison.

To determine the number of Sun Ultra 2 workstations that equal the performance of a 512 processor Cray T3E, a simple mathematical calculation is done with the peak performance of each platform. Note that an analysis of the sustained performances of each platform for executing the same computer code on each platform could lead to different results. It should also be noted that reliability levels of each platform are assumed to be equal.

For this comparison, the peak performance of the Cray T3E is divided by the peak performance of the Sun Ultra 2 workstation (a dual processor machine) to give the ratio of peak performance. The result of this calculation yields the number of Sun Ultra 2 workstations that are required to match the peak performance of the Cray T3E. If the ratio is fractional, the number of machines is rounded up by one.

TABLE III.—NETWORK OF WORKSTATIONS PERFORMANCE  
VERSUS CRAY T3E PERFORMANCE

Platform	Peak performance	Number of processors	Ratio of peak performance	Equivalent number of processors	Equivalent number of machines
Sun Ultra 2	842 Mflops	2	363.42	728	364
Cray T3E	306,000 Mflops	512	1.00	512	1

Table III shows that 364 Sun Ultra 2 workstations are equal in peak performance to the 512 processor Cray T3E. A cost comparison can then be completed with this information.

The SEWP government contract was used to determine the cost of the workstations, file servers, and yearly maintenance prices. The networking and system administration costs are typical costs incurred at NASA Glenn Research Center. The costs assumed for table II are as follows:

- Unix workstation cost = \$22 K each for the Ultra 2 (dual processor, 200-MHz chip, 192-MB memory, 1-GB hard disk)
- Unix workstation network cost = \$1144 per workstation (includes cost of equipment and FTE support)
- Unix workstation server cost = \$30 K each
- One server for every 80 workstations is needed
- Unix workstation maintenance: \$1500 per year per workstation
- Unix workstation system administration costs = 1 system administrator (SA) for 50 Unix workstations. Cost of one SA = \$75 K.

The costs for the Cray T-3E are NASA's costs.

TABLE IV.—NETWORK OF WORKSTATIONS' COST VERSUS CRAY T3E COST

Platform	Equivalent number of machines	Hardware cost	Network cost	Server cost	Maintenance cost	Administration cost	Total cost	Workstation/T3E ratio, percent
Sun Ultra 2	364	\$8,008,000	\$416,416	\$150,000	\$546,000	\$546,000	\$9,666,416	36.36
Cray T3E	1	25,000,000	Included in hardware costs	N/A	864,000	720,000	26,584,000	

Table IV shows that the equivalent in peak performance Sun Ultra 2 workstations costs approximately 36 percent of the Cray T3E. Although this does not meet the CAS Level 1 milestone goal of 25 percent costs, it still shows that workstations are an affordable alternative.

#### REFERENCE

1. DesSureault, T.M.; Roberts, R.G.; and Rothhauser, C.H.: Task 2.1.2.6 Desktop ATM on the Pratt & Whitney Network. Affordable High Performance Computing (AHPC) report, High Performance Computing and Communications Program NASA Glenn Research Center, NASA Cooperative Agreement Notice NCC3-413, May 1998.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 2000	3. REPORT TYPE AND DATES COVERED Technical Memorandum		
4. TITLE AND SUBTITLE  Demonstration of Cost-Effective, High-Performance Computing at Performance and Reliability Levels Equivalent to a 1994 Vector Supercomputer		5. FUNDING NUMBERS  WU-509-10-11-00		
6. AUTHOR(S)  Theresa Babrauckas				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)  National Aeronautics and Space Administration John H. Glenn Research Center at Lewis Field Cleveland, Ohio 44135-3191		8. PERFORMING ORGANIZATION REPORT NUMBER  E-12088		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)  National Aeronautics and Space Administration Washington, DC 20546-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER  NASA TM-2000-209797		
11. SUPPLEMENTARY NOTES  Responsible person, Theresa Babrauckas, organization code 5880, (216) 433-8723.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT  Unclassified - Unlimited Subject Categories: 62 and 01  This publication is available from the NASA Center for AeroSpace Information, (301) 621-0390.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)  The Affordable High Performance Computing (AHPC) project demonstrated that high-performance computing based on a distributed network of computer workstations is a cost-effective alternative to vector supercomputers for running CPU and memory intensive design and analysis tools. The AHPC project created an integrated system called a Network Supercomputer. By connecting computer workstations through a network and utilizing the workstations when they are idle, the resulting distributed-workstation environment has the same performance and reliability levels as the Cray C90 vector supercomputer at less than 25 percent of the C90 cost. In fact, the cost comparison between a Cray C90 Supercomputer and Sun workstations showed that the number of distributed networked workstations equivalent to a C90 costs approximately 8 percent of the C90.				
14. SUBJECT TERMS  Cost reduction; Computer networks; Cray computers; Parallel computers; Workstations; Supercomputers; UNIX operating system			15. NUMBER OF PAGES 13	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT	